arXiv:1605.07371v1 [q-bio.NC] 24 May 2016

Semiparametric energy-based probabilistic models

Jan Humplik* and Gašper Tkačik
Institute of Science and Technology Austria, Klosterneuburg, Austria

Probabilistic models can be defined by an energy function, where the probability of each state is proportional to the exponential of the state's negative energy. This paper considers a generalization of energy-based models in which the probability of a state is proportional to an arbitrary positive, strictly decreasing, and twice differentiable function of the state's energy. The precise shape of the nonlinear map from energies to unnormalized probabilities has to be learned from data together with the parameters of the energy function. As a case study we show that the above generalization of a fully visible Boltzmann machine yields an accurate model of neural activity of retinal ganglion cells. We attribute this success to the model's ability to easily capture distributions whose probabilities span a large dynamic range, a possible consequence of latent variables that globally couple the system. Similar features have recently been observed in many datasets, suggesting that our new method has wide applicability.

I. INTRODUCTION

A probabilistic model over a discrete state space is classified as energy-based if it can be written in the form

$$\begin{aligned} p(\mathbf{s};\ \boldsymbol{\alpha}) &= \frac{e^{-E(\mathbf{s};\ \boldsymbol{\alpha})}}{Z(\boldsymbol{\alpha})}, \\ Z(\boldsymbol{\alpha}) &= \sum_{\mathbf{s}} e^{-E(\mathbf{s};\ \boldsymbol{\alpha})}, \end{aligned} \tag{1}$$

where the *energy* $E(\mathbf{s};\ \boldsymbol{\alpha})$ is a computationally tractable function of the system's configuration $\mathbf{s} = (s_1, s_2, \ldots, s_N)$, $\boldsymbol{\alpha}$ is a set of parameters to be learned from data, and $Z(\boldsymbol{\alpha})$ is a normalization constant also known as the *partition function*. Many methods for learning energy-based models exist (e.g. [1, 2]) which makes them useful in a wide variety of fields, most recently as data analysis tools in neuroscience [3, 4].

A popular way of parametrizing the energy function is by decomposing it into a sum of *potentials* representing interactions among different groups of variables, i.e.

$$\begin{aligned} E(\mathbf{s};\ \boldsymbol{\alpha}) = &\sum_{i=1}^{N} \Phi_i(s_i;\ \boldsymbol{\alpha}_i) + \sum_{i,j=1}^{N} \Phi_{ij}(s_i, s_j;\ \boldsymbol{\alpha}_{ij}) + \\ &+ \sum_{i,j,k=1}^{N} \Phi_{ijk}(s_i, s_j, s_k;\ \boldsymbol{\alpha}_{ijk}) + \ldots. \end{aligned} \tag{2}$$

The resulting models, also termed Gibbs random fields [5], are easy to interpret but, even for moderate N, learning the potential functions from data is intractable unless we can a priori set most of them to zero, or we know how parameters of multiple potentials relate to each other. A common assumption is to consider single- and two-variable potentials only, but many tasks require an efficient parametrization of higher-order interactions.

A powerful way of modeling higher-order dependencies is to assume that they are mediated through hidden variables coupled to the observed system. Many hidden variable models are, however, notoriously hard to learn (e.g. Boltzmann machines [6]), and their distributions over observed variables cannot be represented with tractable energy functions. An important exception is the restricted Boltzmann machine (RBM) [7] which is simple to learn even when the dimension of the data is large, and which has proven effective in many applications [4, 8, 9].

This paper considers a new alternative for modeling higher-order interactions. We generalize any model of the form (1) to

$$\begin{aligned} p(\mathbf{s};\ \boldsymbol{\alpha}, V) &= \frac{e^{-V(E(\mathbf{s};\ \boldsymbol{\alpha}))}}{Z(\boldsymbol{\alpha}, V)}, \\ Z(\boldsymbol{\alpha}, V) &= \sum_{\mathbf{s}} e^{-V(E(\mathbf{s};\ \boldsymbol{\alpha}))}, \end{aligned} \tag{3}$$

where V is an arbitrary strictly increasing and twice differentiable function which needs to be learned from data together with the parameters $\boldsymbol{\alpha}$. While this defines a new energy-based model with an energy function $E'(\mathbf{s};\ \boldsymbol{\alpha}, V) = V(E(\mathbf{s};\ \boldsymbol{\alpha}))$, we will keep referring to $E(\mathbf{s};\ \boldsymbol{\alpha})$ as the energy function. This terminology reflects our interpretation that $E(\mathbf{s};\ \boldsymbol{\alpha})$ should parametrize local interactions between small groups of variables, e.g. low-order terms in Eq. (2), while the function V globally couples the whole system. We will formalize this intuition in Section V. Since setting $V(E) = E$ recovers (1), we will refer to V simply as *the nonlinearity*.

Generalized energy-based models have been previously studied in the physics literature on nonextensive statistical mechanics [10] but, to our knowledge, they have never been considered as data-driven generative models. If $\mathbf{s}$ is a continuous rather than a discrete vector, then models (3) are related to elliptically symmetric distributions [11].

II. NON-PARAMETRIC MODEL OF THE NONLINEARITY

We wish to make as few prior assumptions about the shape of the nonlinearity V as possible. We restrict

* jhumplik@ist.ac.at

ourselves to the class of strictly monotone twice differentiable functions for which V''/V' is square-integrable. It is proved in [12] that any such function can be represented in terms of a square-integrable function W and two constants γ_1 and γ_2 as

$$V(E) = \gamma_1 + \gamma_2 \int_{E_0}^{E} \exp\left(\int_{E_0}^{E'} W(E'')\,\mathrm{d}E''\right)\mathrm{d}E', \quad (4)$$

where E_0 is arbitrary and sets the constants to $\gamma_1 = V(E_0)$, $\gamma_2 = V'(E_0)$. Eq. (4) is a solution to the differential equation $V'' = WV'$, and so W is a measure of the local curvature of V. In particular, V is a linear function on any interval on which $W = 0$.

The advantage of writing the nonlinearity in the form (4) is that we can parametrize it by expanding W in an arbitrary basis without imposing any constraints on the coefficients of the basis vectors. This will allow us to use unconstrained optimization techniques during learning.

We will use piecewise constant functions to parametrize W. Let $[E_0, E_1]$ be an interval containing the range of energies $E(\mathbf{s};\ \boldsymbol{\alpha})$ which we expect to encounter during learning. We divide the interval $[E_0, E_1]$ into Q non-overlapping bins of the same width with indicator functions I_i, i.e. $I_i(E) = 1$ if E is in the ith bin, otherwise $I_i(E) = 0$, and we set $W(E) \equiv W(E;\ \boldsymbol{\beta}) = \sum_{i=1}^{Q} \beta_i I_i(E)$. The integrals in Eq. (4) can be carried out analytically for this choice of W yielding an exact expression for V as a function of $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, as well as for its gradient with respect to these parameters (see Appendix A).

The range $[E_0, E_1]$ and the number of bins Q are meta-parameters which potentially depend on the number of samples in the training set, and which should be chosen by cross-validation.

III. MAXIMUM LIKELIHOOD LEARNING

Any off-the-shelf method for learning energy-based models is suitable for learning the parameters of the nonlinearity and the energy function simultaneously. The nature of these two sets of parameters is, however, expected to be very different, and an algorithm that takes this fact into account explicitly would likely be useful. As a step in this direction, we present an approximation to the likelihood function of the models (3) which can be used to efficiently learn the nonlinearity when the parameters of the energy function are fixed.

A. The nonlinearity enforces a match between the model and the data probability distributions of the energy

Let $\rho(E';\ \boldsymbol{\alpha}) = \sum_{\mathbf{s}} \delta_{E', E(\mathbf{s};\ \boldsymbol{\alpha})}$ count the number of states which map to the same energy E'. The probability distribution of $E(\mathbf{s};\ \boldsymbol{\alpha})$ when $\mathbf{s}$ is distributed according to (3) is

$$\begin{aligned} p(E';\ \boldsymbol{\alpha}, V) &= \sum_{\mathbf{s}} p(\mathbf{s};\ \boldsymbol{\alpha}, V)\delta_{E', E(\mathbf{s};\ \boldsymbol{\alpha})} \\ &= \frac{\rho(E';\ \boldsymbol{\alpha})e^{-V(E')}}{Z(\boldsymbol{\alpha}, V)}. \end{aligned} \quad (5)$$

Given data $\{\mathbf{s}^{(i)}\}_{i=1}^{M}$, let $\hat{p}(E';\ \boldsymbol{\alpha}) = \frac{1}{M}\sum_{i=1}^{M} \delta_{E', E(\mathbf{s}^{(i)};\ \boldsymbol{\alpha})}$ be the data distribution of the energy, and let $\Omega_{\boldsymbol{\alpha}}$ be the image of $E(\mathbf{s};\ \boldsymbol{\alpha})$. The average log-likelihood of the data can be rewritten as

$$\begin{aligned} L(\boldsymbol{\alpha}, V) &= -\log Z(\boldsymbol{\alpha}, V) - \sum_{E' \in \Omega_{\boldsymbol{\alpha}}} \hat{p}(E';\ \boldsymbol{\alpha})V(E') \\ &= -\sum_{E' \in \Omega_{\boldsymbol{\alpha}}} \hat{p}(E';\ \boldsymbol{\alpha}) \log \rho(E';\ \boldsymbol{\alpha}) \\ &\quad + \sum_{E' \in \Omega_{\boldsymbol{\alpha}}} \hat{p}(E';\ \boldsymbol{\alpha}) \log p(E';\ \boldsymbol{\alpha}, V), \end{aligned} \quad (6)$$

where the first line is a standard expression for energy based models, and the second line follows by substituting the logarithm of (5).

Eq. (6) has a simple interpretation. The last term, which is the only one depending on V, is the average log-likelihood of the samples $\{E(\mathbf{s}^{(i)};\ \boldsymbol{\alpha})\}_{i=1}^{M}$ under the model $p(E;\ \boldsymbol{\alpha}, V)$, and so, for any $\boldsymbol{\alpha}$, the purpose of the nonlinearity is to reproduce the data probability distribution of the energy.

Our restriction that V is a twice differentiable increasing function can be seen as a way of regularizing learning. If V was arbitrary, then, for any $\boldsymbol{\alpha}$, an upper bound on the likelihood (6) is attained when $\hat{p}(E;\ \boldsymbol{\alpha}) = p(E;\ \boldsymbol{\alpha}, V)$. According to (5), this can be satisfied with any function V (potentially infinite at some points) such that for all $E \in \Omega_{\boldsymbol{\alpha}}$

$$V(E) = \log \rho(E;\ \boldsymbol{\alpha}) - \log \hat{p}(E;\ \boldsymbol{\alpha}) + \text{const.} \quad (7)$$

Energy functions often assign distinct energies to distinct states in which case the choice (7) leads to a model which exactly reproduces the empirical distribution of data, and hence overfits.

B. Integral approximation of the partition function, and of the likelihood

The partition function (3) can be rewritten as

$$Z(\boldsymbol{\alpha}, V) = \sum_{E' \in \Omega_{\boldsymbol{\alpha}}} \rho(E';\ \boldsymbol{\alpha})e^{-V(E')}. \quad (8)$$

The size of $\Omega_{\boldsymbol{\alpha}}$ is generically exponential in the dimension of $\mathbf{s}$, making the above sum intractable for large systems. To make it tractable, we will use a standard trick from statistical mechanics, and approximate the sum over distinct energies by smoothing $\rho(E;\ \boldsymbol{\alpha})$.

Let $[E_{\min}, E_{\max}]$ be an interval which contains the set Ω_α, and which is divided into K non-overlapping bins $I_i \in$

$[E_{\min}, E_{\max}]$ of width $\Delta = (E_{\max} - E_{\min})/K$. Let E_i be the energy in the middle of the ith bin. The partition function can be approximated as

$$\begin{aligned} Z(\boldsymbol{\alpha}, V) &= \sum_{i=1}^{K} \sum_{E' \in I_i} \rho(E';\ \boldsymbol{\alpha}) e^{-V(E')} \\ &= \sum_{i=1}^{K} e^{-V(E_i)} \sum_{E' \in I_i} \rho(E';\ \boldsymbol{\alpha}) + o(\Delta) \qquad (9) \\ &\equiv \sum_{i=1}^{K} e^{-V(E_i)} \bar{\rho}(E_i;\ \boldsymbol{\alpha}) \Delta + o(\Delta), \end{aligned}$$

where $\bar{\rho}(E_i;\ \boldsymbol{\alpha}) = (1/\Delta) \sum_{E' \in I_i} \rho(E';\ \boldsymbol{\alpha})$ is the *density of states*, and $o(\Delta) \to 0$ as $\Delta \to 0$ since V is differentiable.

The approximation $o(\Delta) \approx 0$ is useful for two reasons. First, in most problems, this approximation becomes very good already when the number of bins K is still much smaller than the size of $\Omega_{\boldsymbol{\alpha}}$. Second, an efficient Monte Carlo method, the Wang and Landau algorithm [13, 14], exists for estimating the density of states for any fixed value of $\boldsymbol{\alpha}$.

The approximation (9) yields the following approximation of the average log-likelihood (6)

$$\begin{aligned} L(\boldsymbol{\alpha}, V) \approx &- \log \left(\sum_{i=1}^{K} e^{-V(E_i)} \bar{\rho}(E_i;\ \boldsymbol{\alpha}) \Delta \right) \\ &- \frac{1}{M} \sum_{i=1}^{M} V(E(\mathbf{s}^{(i)};\ \boldsymbol{\alpha})), \end{aligned} \qquad (10)$$

which can be maximized with respect to the parameters of V using any gradient-based optimization technique (after specifying the metaparameters E_{min}, E_{max}, and K).

Many learning algorithms represent information about the model $p(\mathbf{s};\ \boldsymbol{\alpha}, V)$ as a finite list of samples from the model. This representation is necessarily bad at capturing the low probability regions of $p(E;\ \boldsymbol{\alpha}, V)$. According to (6), this means that any such algorithm is expected to be inefficient for learning V in these regions. On the other hand, the density of states $\bar{\rho}$ estimated with the Wang and Landau algorithm carries information about the low probability regions of $p(E;\ \boldsymbol{\alpha}, V)$ with the same precision as about the high probability regions, and so algorithms based on maximizing (10) should be efficient at learning the nonlinearity.

The approximation (10) cannot be used to learn the parameters $\boldsymbol{\alpha}$ of the energy function, and so the above algorithm has to be supplemented with other techniques.

IV. EXPERIMENTS: MODELING NEURAL ACTIVITY OF RETINAL GANGLION CELLS

A major question in neuroscience is how populations of neurons, rather than individual cells, respond to inputs from the environment. It is well documented that single neuron responses are correlated with each other but the precise structure of the underlying redundancy, and its functional role in information processing is only beginning to be unraveled [3].

The response of a population of N neurons during a short time window can be represented as a binary vector $\mathbf{s} \in \{0, 1\}^N$ by assigning a 1 to every neuron which elicited a spike. We pool all population responses recorded during an experiment, and we ask what probabilistic model would generate these samples. This question was first asked in a seminal paper [15] which showed that, for small networks of retinal neurons ($N < 20$), the fully visible Boltzmann machine, or a *pairwise model*,

$$p(\mathbf{s};\ \boldsymbol{J}) = \frac{1}{Z(\boldsymbol{J})} \exp \left(- \sum_{i,j=1}^{N} J_{ij} s_i s_j \right), \qquad (11)$$

is a good description of the data. Later it was realized that for large networks ($40 < N < 120$) pairwise models cannot accurately capture probability distributions of data statistics which average across the whole population such as the total population activity $K(\mathbf{s}) = \sum_{i=1}^{N} s_i$. This issue was solved by the introduction of the so-called *K-pairwise models* [3],

$$\begin{aligned} p(\mathbf{s};\ \boldsymbol{J}, \boldsymbol{\phi}) = \frac{1}{Z(\boldsymbol{J}, \boldsymbol{\phi})} \exp \Bigg(&- \sum_{i,j=1}^{N} J_{ij} s_i s_j \\ &- \sum_{k=0}^{N} \phi_k \delta_{k, K(\mathbf{s})} \Bigg). \end{aligned} \qquad (12)$$

Here we look at the performance of two additional models. A *semiparametric pairwise model*, which is a generalization (3) of the pairwise model, and a restricted Boltzmann machine with $N/2$ hidden units which is known to be an excellent model of higher order dependencies.

The philosophy of our comparison is not to find a state-of-the art model, but rather to contrast several different models of comparable complexity (measured as the number of parameters) in order to demonstrate that the addition of a simple nonlinearity to an energy-based model can result in a significantly better fit to data.

A. Data

We analyze a simultaneous recording from 160 neurons in a salamander's retina which is presented with 297 repetitions of a 19 second natural movie. The data was collected as part of [3], and it contains a total of $\sim 2.8 \times 10^5$ population responses.

Our goal is to compare the performance, measured as the out-of-sample log-likelihood per sample per neuron, of the above models across several subnetworks of different sizes. To this end, we use our data to construct 48 smaller datasets as follows. We randomly select 40 neurons from the total of 160 as the first dataset. Then we augment this

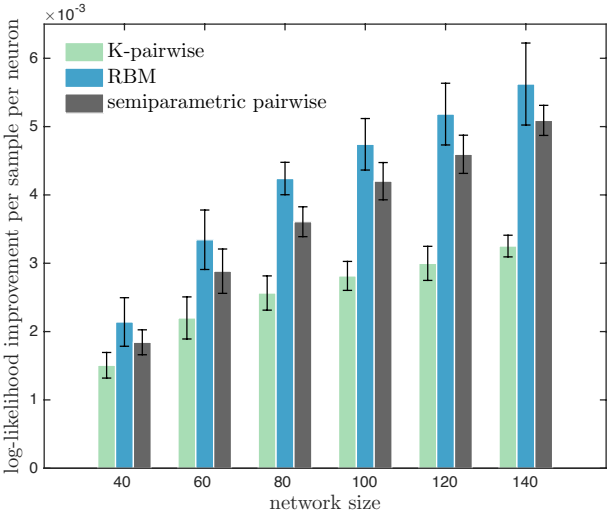

FIG. 1. Out-of-sample log-likelihood improvement per sample per neuron averaged over subnetworks. Errorbars denote variation over subnetworks (one standard deviation), not the accuracy of likelihood estimates. Baseline is the pairwise model.

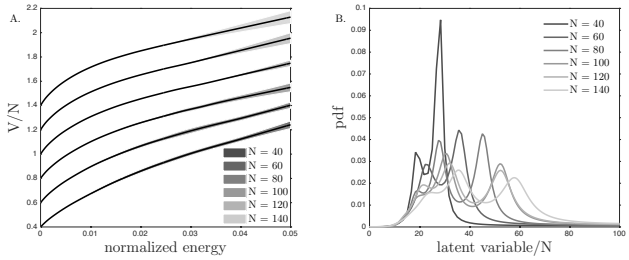

FIG. 2. A. Inferred nonlinearities of the semiparametric pairwise model. The shift along the y-axis is arbitrary, and its only purpose is to increase readability. Black lines are averages over subnetworks. Shaded regions denote variation across subnetworks (one standard deviation). Curves are ordered in increasing order (bottom is $N = 40$, top is $N = 140$). B. Inferred probability densities of the latent variable for one sequence of subnetworks.

dataset with 20 additional neurons to yield the second dataset, and we keep repeating this process until we have a dataset of 140 neurons. This whole process is repeated 8 times, resulting in 8 datasets for each of the 6 different network sizes. Only the datasets with 40 and 60 neurons do not have significant overlaps which is a limitation set by the currently available experiments. For each dataset, we set aside responses corresponding to randomly selected 60 (out of 297) repetitions of the movie, and use these as test data.

B. Results

Our results are summarized in Figure 1. We see that both semiparametric pairwise models and RBMs with $N/2$ hidden units substantially outperform the previously considered K-pairwise models. In particular, the improvement increases as we go to larger networks. RBMs are consistently better than semiparametric pairwise models, but, interestingly, this gap does not seem to scale with the network size.

The inferred nonlinearity V for semiparametric pairwise models does not vary much across subnetworks of the same size. This is not surprising for large networks since there is a substantial overlap between the datasets, but it is nontrivial for smaller networks. The average inferred nonlinearities are shown in Figure 2A. Semiparametric pairwise models have one free parameter since the transformation $\boldsymbol{J} \to \kappa \boldsymbol{J}, V(E) \to V(E/\kappa)$ does not change their probability distribution. Therefore, to make the comparison of V across different datasets fair, we fix κ by requiring $\sum_{i,j} J_{ij}^2 = 1$. Furthermore, the nonlinearities in Figure 2A are normalized by N because we expect the probabilities in a system of dimension N to scale as e^{-N}. The inferred nonlinearities are approximately concave, and their curvature increases with the network size.

C. Training the models

Training was done by a variation of Persistent Contrastive Divergence [1] which performs an approximate gradient ascent on the log-likelihood of any energy-based model (1). Given an initial guess of the parameters $\boldsymbol{\alpha}_0$, and a list of M_s samples drawn from $p(\mathbf{s};\ \boldsymbol{\alpha}_0)$, the algorithm can be summarized as

$$\begin{aligned}&\text{for } t := 1 \text{ to } L\\ &\quad \boldsymbol{\alpha}_t = \boldsymbol{\alpha}_{t-1} + \eta(\mathbf{E}[\nabla_{\boldsymbol{\alpha}} E(\mathbf{s};\ \boldsymbol{\alpha}_{t-1})]_{\text{samples}_{t-1}}\\ &\qquad\qquad -\mathbf{E}[\nabla_{\boldsymbol{\alpha}} E(\mathbf{s};\ \boldsymbol{\alpha}_{t-1})]_{\text{data}})\\ &\quad \text{samples}_t = \text{GIBBS}^n(\text{samples}_{t-1}, \boldsymbol{\alpha}_t)\end{aligned}$$

where L is the number of iterations, η is the learning rate, $\mathbf{E}[\cdot]_{\text{list}}$ denotes an average over the list of states, and GIBBS^n represents n applications (i.e. n neurons are switched) of the Gibbs sampling transition operator.

In the case of semiparametric pairwise models, Persistent Contrastive Divergence was used as a part of an alternating maximization algorithm in which we learned the nonlinearity by maximizing the approximate likelihood (10) while keeping the couplings $\boldsymbol{J}$ fixed, and then learned the couplings using Persistent Contrastive Divergence with the nonlinearity fixed. Details of the learning algorithms for all models are described in Appendix B.

The most interesting metaparameter is the number of bins Q necessary to model the nonlinearity V. We settled on $Q = 12$ but we observed that decreasing it to $Q = 6$ would not significantly change the training likelihood. However, $Q = 12$ yielded more consistent nonlinearities over different subgroups.

D. Estimating likelihood and controlling for overfitting

Estimating likelihood of our data is easy because the state $\mathbf{s}_0 = (0, \ldots, 0)$ occurs with high probability (~ 0.2), and all the inferred models retain this property. Therefore, for any energy-based model, we estimated $p(\mathbf{s}_0;\ \boldsymbol{\alpha})$ by drawing 3×10^6 samples using Gibbs sampling, calculated the partition function as $Z(\boldsymbol{\alpha}) = \exp(-E(\mathbf{s}_0);\ \boldsymbol{\alpha})/p(\mathbf{s}_0;\ \boldsymbol{\alpha})$, and used it to calculate the likelihood.

The models do not have any explicit regularization. We tried to add a smoothed version of L1 regularization on coupling matrices but we did not see any improvement in generalization using a cross-validation on one of the training datasets. Certain amount of regularization is due to sampling noise in the estimates of likelihood gradients, helping us to avoid overfitting.

V. WHY DOES IT WORK?

Perhaps surprisingly, the addition of a simple nonlinearity to the pairwise energy function $E(\mathbf{s};\ \boldsymbol{J}) = \sum_{i,j=1}^{N} J_{ij} s_i s_j$ significantly improves the fit to data. Here we give heuristic arguments that this should be expected whenever the underlying system is globally coupled.

A. Many large complex systems are globally coupled

Let $p_N(\mathbf{s})$ be a sequence of positive probabilistic models ($\mathbf{s}$ is of dimension N), and suppose that $p_N(\mathbf{s})$ can be (asymptotically) factorized into subsystems statistically independent of each other whose number is proportional to N. Then $(1/N) \log p_N(\mathbf{s})$ is an average of independent random variables, and we expect its standard deviation $\sigma(\log p_N(\mathbf{s}))/N$ to vanish in the $N \to \infty$ limit. Alternatively, if $\sigma(\log p_N(\mathbf{s}))/N \not\to 0$, then the system cannot be decomposed into independent subsystems, and there must be some mechanism globally coupling the system together.

It has been argued that many natural systems [16] including luminance in natural images [17], amino acid sequences of proteins [18], and neural activity such as the one studied in Sec. IV [19, 20] belong to the class of models whose log-probabilities per dimension have large variance even though their dimensionality is big. Therefore, models of such systems should reflect the prior expectation that there is a mechanism which couples the whole system together. In our case, this mechanism is the nonlinearity V.

B. Mapping the nonlinearity to a latent variable

Recent work attributes the strong coupling observed in many systems to the presence of latent variables ([21, 22]). We can rewrite the model (3) in terms of a latent variable considered in [21] if we assume that $\exp(-V(E))$ is a totally monotone function, i.e. that it is continuous for $E \geq 0$ (we assume, without loss of generality, that $E(\mathbf{s};\ \boldsymbol{\alpha}) \geq 0$), infinitely differentiable for $E > 0$, and that $(-1)^n d^n \exp(-V(E))/dE^n \geq 0$ for $n \geq 0$. Bernstein's theorem [23] then asserts that we can rewrite the model (3) as

$$\begin{aligned} \frac{e^{-V(E(\mathbf{s};\ \boldsymbol{\alpha}))}}{Z(\boldsymbol{\alpha}, V)} &= \int_0^\infty q(h) \frac{e^{-hE(\mathbf{s};\ \boldsymbol{\alpha})}}{Z(h;\ \boldsymbol{\alpha})}\, \mathrm{d}h, \\ Z(h;\ \boldsymbol{\alpha}) &= \sum_{\mathbf{s}} e^{-hE(\mathbf{s};\ \boldsymbol{\alpha})}, \end{aligned} \tag{13}$$

where $q(h)$ is a probability density (possibly containing delta functions). Suppose that the energy function has the form (2). Then we can interpret (13) as a latent variable h being coupled to every group of interacting variables, and hence inducing a coupling between the whole system whose strength depends on the size of the fluctuations of h.

While the class of positive, twice differentiable, and decreasing functions that we consider is more general than the class of totally monotone functions, we can find the maximum likelihood densities $q(h)$ which correspond to the pairwise energy functions inferred in Section IV using the semiparametric pairwise model. We model $q(h)$ as a histogram, and maximize the likelihood under the model (13) by estimating $Z(h;\ \boldsymbol{J})$ using the approximation (9). The maximum likelihood densities are shown in Figure 2B for one particular sequence of networks of increasing size. The units of the latent variables are arbitrary, and set by the scale of $\boldsymbol{J}$ which we normalize so that $\sum_{i,j} J_{ij}^2 = 1$. The bimodal structure of the latent variables is observed across all datasets. We do not observe a significant decrease in likelihood by replacing the nonlinearity V with the integral form (13). Therefore, at least for the data in Sec. IV, the nonlinearity can be interpreted as a latent variable globally coupling the system.

C. Asymptotic form of the nonlinearity

Suppose that the true system $\hat{p}_N(\mathbf{s})$ which we want to model with (3) satisfies $\sigma(\log p_N(\mathbf{s}))/N \not\to 0$. Then we also expect that energy functions $E_N(\mathbf{s};\ \boldsymbol{\alpha})$ which accurately model the system satisfy $\sigma(E_N(\mathbf{s};\ \boldsymbol{\alpha}))/N \not\to 0$. In the limit $N \to \infty$, the approximation of the likelihood (10) becomes exact when V is differentiable (Δ has to be appropriately scaled with N), and arguments which led to (7) can be reformulated in this limit to yield

$$V_N(E) = \log \bar{\rho}_N(E;\ \boldsymbol{\alpha}) - \log \hat{\bar{p}}_N(E;\ \boldsymbol{\alpha}) + \text{const.}, \tag{14}$$

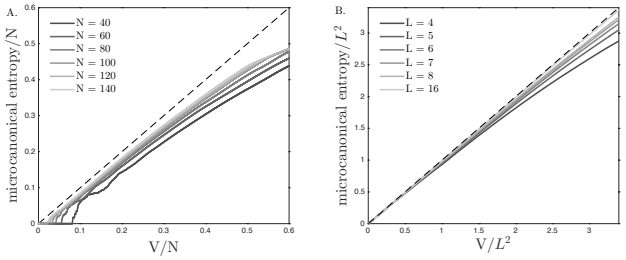

FIG. 3. A. Plot of the inferred nonlinearity vs the microcanonical entropy for the semiparametric pairwise model. B. The same for the elliptically symmetrical model of luminance distribution in natural images.

where $\bar{\rho}_N(E;\ \boldsymbol{\alpha})$ is the density of states, and $\hat{\bar{p}}_N(E;\ \boldsymbol{\alpha})$ is now the probability density of $E_N(\mathbf{s};\ \boldsymbol{\alpha})$ under the true model. In statistical mechanics, the first term $\log \bar{\rho}_N(E;\ \boldsymbol{\alpha})$ is termed the *microcanonical entropy*, and is expected to scale linearly with N. On the other hand, because $\sigma(E_N(\mathbf{s};\ \boldsymbol{\alpha}))/N \nrightarrow 0$, we expect the second term to scale at most as $\log N$. Thus we make the prediction that if the underlying system cannot be decomposed into independent subsystems, then the maximum likelihood nonlinearity satisfies $V(E) \approx \log \bar{\rho}(E;\ \boldsymbol{\alpha})$ up to an arbitrary constant.

In Figure 3A we show a scatter plot of the inferred nonlinearity for one sequence of subnetworks in Sec. IV vs the microcanonical entropy estimated using the Wang and Landau algorithm. While the convergence is slow, the plot suggests that these functions approach each other as the network size increases. To demonstrate this prediction on yet another system, we used the approach in [11] to fit the semiparametric pairwise model with $\mathbf{s} \in \mathbb{R}^{L^2}$ to $L \times L$ patches of pixel log-luminances in natural scenes from the database [24]. This model has an analytically tractable density of states which makes the inference simple. Figure 3B shows the relationship between the inferred nonlinearity and the microcanonical entropy for collections of patches which increase in size, confirming our prediction that the nonlinearity should be given by the microcanonical entropy.

VI. CONCLUSIONS

We presented a tractable extension of any energy-based model which can be interpreted as augmenting the original model with a latent variable. As demonstrated on the retinal activity data, this extension can yield a substantially better fit to data even though the number of additional parameters is negligible compared to the number of parameters of the original model. In light of our results, we hypothesize that combing a nonlinearity with the energy function of a restricted Boltzmann machine might yield a model of retinal activity which is not only accurate, but also simple as measured by the number of parameters. Simplicity is an important factor in neuroscience because of experimental limitations on the number of samples. We plan to pursue this hypothesis in future work.

Our models are expected to be useful whenever the underlying system cannot be decomposed into independent components. This phenomenon has been observed in many natural systems, and the origins of this global coupling, and especially its analogy to physical systems at critical points, have been hotly debated. Our models effectively incorporate the prior expectations of a global coupling in a simple nonlinearity, making them superior to models based on Gibbs random fields which might need a large number of parameters to capture the same dependency structure.

ACKNOWLEDGMENTS

We thank David Schwab, Elad Schneidman, and William Bialek for helpful discussions. This work was supported in part by HFSP Program Grant RGP0065/2012 and Austrian Science Fund (FWF) grant P25651.

Appendix A: Exact expression for the nonlinarity and its gradient

Let $W(E;\ \boldsymbol{\beta}) = \sum_{i=1}^{Q} \beta_i I_i(E)$ where I_i are the indicator functions defined in the main text. For $E < E_0$, we have $V(E;\ \boldsymbol{\gamma}, \boldsymbol{\beta}) = \gamma_1 + \gamma_2(E - E_0)$. For $E > E_0$ we have $V(E;\ \boldsymbol{\gamma}, \boldsymbol{\beta}) = \gamma_1 + \gamma_2 f(E;\ \boldsymbol{\beta})$, where

$$\begin{aligned} f(E;\ \boldsymbol{\beta}) &= \int_{E_0}^{E} \exp\left(\int_{E_0}^{E'} W(E'';\ \boldsymbol{\beta})\,\mathrm{d}E''\right)\mathrm{d}E' \\ &= \sum_{i=1}^{[E]-1} \exp\left(\Delta \sum_{j=1}^{i-1} \beta_j\right) \frac{\exp(\Delta\beta_i) - 1}{\beta_i} + \exp\left(\Delta \sum_{j=1}^{[E]-1} \beta_j\right) \frac{\exp(\beta_{[E]}(E - ([E]-1)\Delta)) - 1}{\beta_{[E]}}. \end{aligned} \tag{A1}$$

We define $[E]$ as the number of the bin that contains E. If $E > E_1$, then we define $[E] = Q + 1$, and $\beta_{Q+1} = 0$.

The gradient is

$$\frac{\partial V(E;\ \boldsymbol{\gamma}, \boldsymbol{\beta})}{\partial \gamma_1} = 1, \tag{A2}$$

$$\frac{\partial V(E;\,\boldsymbol{\gamma},\boldsymbol{\beta})}{\partial \gamma_2} = f(E;\,\boldsymbol{\beta}), \tag{A3}$$

$$\frac{\partial V(E;\,\boldsymbol{\gamma},\boldsymbol{\beta})}{\partial \beta_k} = \gamma_2 \frac{f(E;\,\boldsymbol{\beta})}{\partial \beta_k}. \tag{A4}$$

If $k > [E]$, then

$$\frac{\partial f(E;\,\boldsymbol{\beta})}{\partial \beta_k} = 0. \tag{A5}$$

If $k = [E]$, then

$$\frac{\partial f(E;\,\boldsymbol{\beta})}{\partial \beta_k} = \exp\left(\Delta \sum_{j=1}^{[E]-1} \beta_j\right) \frac{\exp(\Delta\beta_{[E]})\Delta\beta_{[E]} - \exp(\Delta\beta_{[E]}) + 1}{\beta_{[E]}^2}. \tag{A6}$$

If $k < [E]$, then

$$\begin{aligned}\frac{\partial f(E;\,\boldsymbol{\beta})}{\partial \beta_k} = \exp\left(\Delta \sum_{j=1}^{k-1} \beta_j\right) \frac{\exp(\Delta\beta_k)\Delta\beta_k - \exp(\Delta\beta_k) + 1}{\beta_k^2} + \Delta \sum_{i=k+1}^{[E]-1} \exp\left(\Delta \sum_{j=1}^{i-1} \beta_j\right) \frac{\exp(\Delta\beta_i) - 1}{\beta_i} \\ + \Delta \exp\left(\Delta \sum_{j=1}^{[E]-1} \beta_j\right) \frac{\exp(\beta_{[E]})(E - ([E]-1)\Delta) - 1}{\beta_{[E]}}.\end{aligned} \tag{A7}$$

Appendix B: Details of training

1. Pairwise models, K-pairwise models, and RBMs

Pairwise models, K-pairwise models, and RBMs were all trained using Persistent Contrastive Divergence with $\eta = 2$, $n = 2N$, and with initial parameters drawn from a normal distribution with 0 mean and 0.1 standard deviation. We iterated the algorithm three times, first with $L = 1000, M_s = 3 \times 10^4$, then with $L = 1000, M_s = 3 \times 10^5$, and finally with $L = 500, M_s = 3 \times 10^6$ (the last step had $L = 1000$ for K-pairwise models, and also for RBMs when $N = 140$).

2. Semiparametric pairwise models

We initialized the coupling matrix $\boldsymbol{J}_0$ as the one learned using a pairwise model. The E_0 and E_1 metaparameters of V were set to the minimum and maximum energy $E(\mathbf{s};\,\boldsymbol{J}_0) = \sum_{ij} J_{0ij} s_i s_j$ observed in the training set. We set $Q = 6$, and we initialized the parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ of the nonlinearity by maximizing the approximate likelihood with $\boldsymbol{J} = \boldsymbol{J}_0$ fixed. The metaparameters $E_{\min}$ and $E_{\max}$ for the approximate likelihood were set to the minimum, and twice the maximum of $E(\mathbf{s};\,\boldsymbol{J})$) over the training set. K was set between 2000 and 3000. The density of states was estimated with a variation of the algorithm described in [14] with accuracy $F_{\text{final}} = 10^{-8}$.

Starting with these initial parameters, we ran two iterations of Persistent Contrastive Divergence with $n = 2N$, simultaneously learning the coupling matrix and the nonlinearity. The first iteration had $L = 1000, M_s = 3 \times 10^4$, and the second one $L = 1000, M_s = 3 \times 10^5$. In order for the learning to be stable, we had to choose different learning rates for the coupling matrix ($\eta = 2$), and for $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ ($\eta = 10^{-4}$).

For the last step, we adjusted the metaparameters of the nonlinearity so that E_0 and E_1 are the minimum and maximum observed energies with the current $\boldsymbol{J}$, and $Q = 12$. We maximized the approximate likelihood to infer the nonlinearity with these new metaparameters. Then we fixed V, and ran a Persistent Contrastive Divergence ($\eta = 2, n = 2N, L = 500, M_s = 3 \times 10^6$) learning $\boldsymbol{J}$. Finally we maximized the approximate likelihood with fixed $\boldsymbol{J}$ to get the final nonlinearity.

[1] Tijmen Tieleman. Training restricted boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th International Conference on*

Machine Learning, ICML '08, pages 1064–1071, New York, NY, USA, 2008. ACM.

[2] Jascha Sohl-Dickstein, Peter B. Battaglino, and Michael R. DeWeese. New method for parameter estimation in probabilistic models: Minimum probability flow. *Phys. Rev. Lett.*, 107:220601, Nov 2011.

[3] Gašper Tkačik, Olivier Marre, Dario Amodei, Elad Schneidman, William Bialek, and Michael J. Berry, II. Searching for collective behavior in a large network of sensory neurons. *PLoS Comput Biol*, 10(1):e1003408, 01 2014.

[4] Urs Köster, Jascha Sohl-Dickstein, Charles M. Gray, and Bruno A. Olshausen. Modeling higher-order correlations within cortical microcolumns. *PLoS Comput Biol*, 10(7):e1003684, 07 2014.

[5] R. Kindermann and J. L. Snell. *Markov Random Fields and Their Applications*. AMS books online. American Mathematical Society, 1980.

[6] G. E. Hinton and T. J. Sejnowski. Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1. chapter Learning and Relearning in Boltzmann Machines, pages 282–317. MIT Press, Cambridge, MA, USA, 1986.

[7] P. Smolensky. Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1. chapter Information Processing in Dynamical Systems: Foundations of Harmony Theory, pages 194–281. MIT Press, Cambridge, MA, USA, 1986.

[8] Geoffrey E Hinton and Ruslan R Salakhutdinov. Replicated softmax: an undirected topic model. In *Advances in neural information processing systems*, pages 1607–1614, 2009.

[9] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. Restricted boltzmann machines for collaborative filtering. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, pages 791–798, New York, NY, USA, 2007. ACM.

[10] R. Hanel and S. Thurner. When do generalized entropies apply? How phase space volume determines entropy. *EPL (Europhysics Letters)*, 96(5):50003, 2011.

[11] Siwei Lyu and Eero P Simoncelli. Reducing statistical dependencies in natural signals using radial gaussianization. In *Advances in neural information processing systems*, pages 1009–1016, 2009.

[12] J. O. Ramsay. Estimating smooth monotone functions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(2):365–375, 1998.

[13] Fugao Wang and D. P. Landau. Efficient, multiple-range random walk algorithm to calculate the density of states. *Phys. Rev. Lett.*, 86:2050–2053, Mar 2001.

[14] R. E. Belardinelli and V. D. Pereyra. Fast algorithm to calculate density of states. *Phys. Rev. E*, 75:046701, Apr 2007.

[15] Elad Schneidman, Michael J. Berry, II, Ronen Segev, and William Bialek. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440(7087):1007–12, 2006.

[16] Thierry Mora and William Bialek. Are Biological Systems Poised at Criticality? *Journal of Statistical Physics*, 144(2):268–302, 2011.

[17] Greg J. Stephens, Thierry Mora, Gašper Tkačik, and William Bialek. Statistical Thermodynamics of Natural Images. *Physical Review Letters*, 110(1):018701, 2013.

[18] Thierry Mora, Aleksandra M. Walczak, William Bialek, and Curtis G. Callan. Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences of the United States of America*, 107(12):5405–5410, 2010.

[19] Gašper Tkačik, Thierry Mora, Olivier Marre, Dario Amodei, Stephanie E. Palmer, Michael J. Berry, II, and William Bialek. Thermodynamics and signatures of criticality in a network of neurons. *Proceedings of the National Academy of Sciences*, 112(37):11508–11513, 2015.

[20] Thierry Mora, Stéphane Deny, and Olivier Marre. Dynamical criticality in the collective activity of a population of retinal neurons. *Phys. Rev. Lett.*, 114(7):078105, 2015.

[21] David J. Schwab, Ilya Nemenman, and Pankaj Mehta. Zipf's Law and Criticality in Multivariate Data without Fine-Tuning. *Phys. Rev. Lett.*, 113(6):068102, 2014.

[22] Laurence Aitchison, Nicola Corradi, and Peter E Latham. Zipf's law arises naturally in structured, high-dimensional data. *arXiv preprint arXiv:1407.7135*, 2014.

[23] D.V. Widder. *The Laplace transform*. Princeton mathematical series. Princeton university press, 1946.

[24] Gašper Tkačik, Patrick Garrigan, Charles Ratliff, Grega Milčinski, Jennifer M. Klein, Lucia H. Seyfarth, Peter Sterling, David H. Brainard, and Vijay Balasubramanian. Natural images from the birthplace of the human eye. *PLoS ONE*, 6(6):e20409, 06 2011.